

Bayesian time series analysis of terrestrial impact cratering

C.A.L. Bailer-Jones*

Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

Submitted 13 March 2011. Revised 11 April 2011 and 17 May 2011. Accepted 20 May 2011. Minor typos corrected 15 June 2011. Minor correction (erratum) 16 October 2011

ABSTRACT

Giant impacts by comets and asteroids have probably had an important influence on terrestrial biological evolution. We know of around 180 high velocity impact craters on the Earth with ages up to 2400 Myr and diameters up to 300 km. Some studies have identified a periodicity in their age distribution, with periods ranging from 13 to 50 Myr. It has further been claimed that such periods may be causally linked to a periodic motion of the solar system through the Galactic plane. However, many of these studies suffer from methodological problems, for example misinterpretation of p-values, overestimation of significance in the periodogram or a failure to consider plausible alternative models. Here I develop a Bayesian method for this problem in which impacts are treated as a stochastic phenomenon. Models for the time variation of the impact probability are defined and the evidence for them in the geological record is compared using Bayes factors. This probabilistic approach obviates the need for ad hoc statistics, and also makes explicit use of the age uncertainties. I find strong evidence for a monotonic decrease in the recorded impact rate going back in time over the past 250 Myr for craters larger than 5 km. The same is found for the past 150 Myr when craters with upper age limits are included. This is consistent with a crater preservation/discovery bias modulating an otherwise constant impact rate. The set of craters larger than 35 km (so less affected by erosion and infilling) and younger than 400 Myr are best explained by a constant impact probability model. A periodic variation in the cratering rate is strongly disfavoured in all data sets. There is also no evidence for a periodicity superimposed on a constant rate or trend, although this more complex signal would be harder to distinguish.

Key words: methods: data analysis, statistical – Earth – meteorites, meteors, meteoroids – planets and satellites: surfaces

1 INTRODUCTION

About 180 terrestrial impact craters are known. The high velocity of the impact means that even relatively small comets or asteroids produce large craters. The meteor which crated the 1.2 km diameter Barringer crater in Arizona, for example, was probably only 50m across. Since the discovery of evidence for a large impact 65 Myr ago at the geological boundary between the Cretaceous and Tertiary periods (the K-T boundary) (Alvarez et al. 1980) and its implication in the mass extinction event at that time (including the demise of the dinosaurs), it has become clear that bolide impacts have had a significant impact on the evolution of life.

The large impactors are believed to be either asteroids from the main asteroid belt, or comets from the Oort cloud (Shoemaker 1983). The multi-body dynamics involved in putting these on a collision course with the Earth implies that cratering is a random phenomenon, but the rate of impacts is not necessarily constant in time. It has been suggested that gravitational perturbations of the Oort

cloud due to the Galactic tide, passages of the solar system near to molecular clouds, or an unseen solar companion, may send large numbers of comets into the inner solar system as a comet shower, increasing the impact rate (Davis et al. 1984, Torbett & Smoluchowski 1984, Rampino & Stothers 1984, Napier 1998, Wickramasinghe & Napier 2008, Gardner et al. 2011). Simple dynamical calculations indicate that the Sun oscillates vertically about the Galactic midplane with a period of 52–74 Myr, depending on the mass density assumed (Bahcall & Bahcall 1985, Shuter & Klatt 1986, Stothers 1998). In parallel to this, several studies claim to have found evidence for a temporal periodicity in the impact cratering record over the past few hundred million years, with numerous periods ranging from 13 to 50 Myr having been identified (Alvarez & Muller 1984, Rampino & Stothers 1984, Grieve et al. 1985, Montanari et al. 1998, Napier 1998, Yabushita 2004, Chang & Moon 2005, Napier 2006). Some authors make a causal link, suggesting that each midplane crossing increases the impact rate. It has also been suggested that periodicities in cratering may be associated with alleged periodicities in mass extinctions or in biodiversity variations, although there is little evidence linking any mass extinc-

* Email: calj@mpia.de

tion apart from the K-T one to a giant impact (Alvarez 2003, Hal-lam 2004; see Bailer-Jones 2009 for a more general review of ex-traterrestrial influence on terrestrial climate and biodiversity.) Other studies of the crater record conclude there to be insufficient evi-dence for periodicity (e.g. Grieve 1991, Grieve & Pesonen 1996, Yabushita 1996, Jetsu & Pelt 2000).

While these are a priori reasonable suggestions worthy of fur-ther analysis, many of the studies claiming to have identified peri-ods are compromised by problems with their methodology. Typical problems are misinterpreting p-values, overestimating the signifi-cance of periodogram peaks, or failing to consider a sufficient set of models. Possibly because the data comprise only crater ages (with no attached magnitudes like more familiar time series), many stud-ies have developed ad hoc statistics to look for periods, many of which have poorly explored statistical properties. Identifying “pe-riods” is relatively easy – any time series can be expressed as a sum of Fourier terms; clusters of points can always be found – but properly assessing significance is harder. A common mistake is to interpret evidence *against* a null hypothesis of “random data” as evidence *for* some periodic model, neglecting to realise that both may be inferior to a plausible third alternative. Although these are known limitations of frequentist hypothesis testing which have been discussed extensively (e.g. Berger & Sellke 1987, Kass & Raftery 1996, Marden 2000, Berger 2003, Jaynes 2003, MacKay 2003, Christensen 2005), this seems not to have discouraged their use.

The aim of this article is to analyse the crater time series with well-motivated statistical methods. One of the key features is to write down explicit models for the impact phenomenon. A second feature is that I consider this phenomenon to be a *stochastic* pro-cess: Rather than expecting the impact events to follow a determin-istic pattern, I model the time variation of the impact probability. This better accommodates the astrophysical and geological con-texts (e.g. smooth variations in the torques on the Oort cloud, or slow erosion of craters). By using the Bayesian framework to anal-yse the data, we can properly calculate the evidence for the different models and compare them on an equal footing. The critical aspect is that we must compare evidence for the entire model (the average likelihood over the model parameter space), rather than compar-ing the tuned maximum likelihood fit (which generally favours the more complex model).

Craters are difficult to date, and some have very large age un-certainties (Grieve 1991, Deutsch & Schärer 1994). There has been a lot of agonizing in the literature about how to deal with age un-certainties, and some studies may have biased their conclusions by re-moving “poor” data. Contrary to claims in the literature, including craters with large uncertainties is not a problem for model compar-ison *provided* we use these uncertainties appropriately. The method developed in this paper does this, and it can also include craters which only have upper age limits.

The impact crater record is certainly incomplete. Of most rel-evance is the preservation bias: on account of erosion and infilling, older, smaller craters are less likely to survive or to be discovered. The oldest known crater is about 2400 Myr old, but there is an ob-vious paucity of craters older than about 700 Myr (only 14 of 176 older than this). For this reason I focus my analysis on craters larger than 5 km in diameter which have ages (or upper age limits) be-low 250 Myr, although I also extend the analysis back to 400 Myr BP (before present). Rather than attempting to “de-bias” the data, I model the data as they are, so strictly I am modelling not the orig-inal impact history (which we do not observe), but the combined impact/preservation history.

After outlining the data in section 2, I describe the method in section 3. This is tested and demonstrated on simulated data in sec-tion 4. The results are presented in section 5, followed by a discus-sion of these and the method and a comparison with other studies in section 6. I summarize and conclude in section 7. I will say very little about the wider topics of the comet and asteroid population, impact effects, crater geology and aging etc. These are reviewed by Shoemaker (1983), Grieve (1991), Deutsch & Schärer (1994) and Grieve & Pesonen (1996), amongst others.

2 IMPACT CRATER DATA

The data are taken from the *Earth Impact Database* (EID), a com-pilation from the literature maintained by the Planetary and Space Science Centre at the University of New Brunswick.¹ This has been the source of data for many previously published studies, and has been continuously expanded as new craters have been discovered, and information revised as improved age or size estimates obtained. As of 30 September 2010 this listed 176 craters. My study fo-cuses primarily on craters younger than 250 Myr with diameters $d > 5$ km. 59 craters fulfil these criteria and are listed in Table 1. 42 of these have an age and age uncertainty in the EID. These uncer-tainties are interpreted as 1σ Gaussian uncertainties (see section 3.2 for why this is so). The data on the remaining 17 craters fall into three groups:

- 13 craters have only upper limits to their ages. These can be included in the analysis consistently, as outlined in section 3.5.
- Two craters, *Bosumtwi* and *Haughton*, have no age uncertainty reported in the EID. Craters younger than 50 Myr have dating errors ranging from 0.1 to 20 Myr (or 0.3% to 60%), so it is difficult to estimate an appropriate uncertainty. I rather arbitrarily assign 10% uncertainties to the ages of these two craters.
- Two craters, *Avak* and *Jebel Waqf as Suwwan*, have age ranges (3–95 Myr and 37–56 Myr respectively) rather than estimates with uncertainties in the EID. I consider the range as a 90% confidence interval of a Gaussian distribution with mean equal to the average of the limits (so the standard deviation is 0.30 times the range).

Crater diameters are notoriously difficult to measure. No uncer-tainties are listed in the EID, but the uncertainty can be a factor of two or more for buried craters (e.g. Grieve 1991). On account of erosion and infilling, the older or smaller a crater is, the less likely it is to be preserved or discovered, resulting in increasing incompleteness with look back time. For this reason – and following several other studies – I limit my analysis to craters larger than 5 km in diame-ter. Our knowledge of geological processes suggests that for times back to 150–250 Myr, most craters larger than 5 km should be rea-sonably well preserved (although this is something I will test). Note that I only use the diameters to select the sample; they are not used in the actual analysis.

We certainly have not identified all impacts: the continents have not been explored equally thoroughly, and craters from ma-rine impacts are rarer and harder to identify. But provided these selection effects are time independent (and size independent above 5 km), they do not introduce any relevant biases.

The continents have also moved. 250 Myr, ago at the Permian–Triassic boundary, there was a larger concentration of land mass towards southern latitudes than there is now. Given that asteroid

¹ <http://www.passc.net/EarthImpactDatabase/>

[h]

Table 1. The 59 craters in the Earth Impact Database with diameters greater than or equal to 5 km and ages or age upper limits below 250 Myr

Name	age Myr	$\sigma(\text{age})$ Myr	diameter km
Araguainha	244.4	3.25	40
Avak	49.0	28.0	12
Beyenchime-Salaatin	40	20	8
Bigach	5	3	8
Boltysh	65.17	0.64	24
Bosumtwi	1.07	0.107	10.5
Carswell	115	10	39
Chesapeake Bay	35.3	0.1	90
Chicxulub	64.98	0.05	170
Chiylí	46	7	5.5
Chukcha	< 70		6
Cloud Creek	190	30	7
Connolly Basin	< 60		9
Deep Bay	99	4	13
Dellen	89	2.7	19
Eagle Butte	< 65		10
El'gygytgyn	3.5	0.5	18
Goat Paddock	< 50		5.1
Gosses Bluff	142.5	0.8	22
Haughton	39	3.9	23
Jebel Waqf as Suwwan	46.5	5.8	5.5
Kamensk	49	0.2	25
Kara	70.3	2.2	65
Kara-Kul	< 5		52
Karla	5	1	10
Kentland	< 97		13
Kursk	250	80	6
Lappajrvi	73.3	5.3	23
Logancha	40	20	20
Logoisk	42.3	1.1	15
Manicouagan	214	1	100
Manson	74.1	0.1	35
Maple Creek	< 75		6
Marquez	58	2	12.7
Mien	121	2.3	9
Mistastin	36.4	4	28
Mjlnir	142	2.6	40
Montagnais	50.5	0.76	45
Morokweng	145	0.8	70
Oasis	< 120		18
Obolon'	169	7	20
Popigai	35.7	0.2	100
Puchezh-Katunki	167	3	80
Ragozinka	46	3	9
Red Wing	200	25	9
Ries	15.1	0.1	24
Rochechouart	214	8	23
Saint Martin	220	32	40
Sierra Madera	< 100		13
Steen River	91	7	25
Tin Bider	< 70		6
Tookoonooka	128	5	55
Upheaval Dome	< 170		10
Vargeao Dome	< 70		12
Vista Alegre	< 65		9.5
Wanapitei	37.2	1.2	7.5
Wells Creek	200	100	12
Wetumpka	81	1.5	6.5
Zhamanshin	0.9	0.1	14

impactors are not distributed isotropically – their orbits are concentrated in the ecliptic – we may expect this continental drift to introduce a time variation in the impact rate. However, calculations by Le Feuvre & Wieczorek (2008) show that the impact probability actually has only a very weak latitudinal dependence, being only 4% lower at the poles than at the equator. The net dependence will be even lower when (isotropic) comets are included. Thus continental drift produces a very minor bias which I ignore.

Fig. 1 shows the ages and diameters of the 46 craters in Table 1 which have age estimates (rather than upper limits). Fig. 2 show the ages uncertainties on these points and includes the 13 craters with upper limits on their ages.

I perform the analysis on several different data sets over three ages ranges. Craters with age estimates and uncertainties (σ_t) originating from the EID form the *basic* data sets. Adding to these the craters for which I have assigned ages/age uncertainties forms the *extended* data sets. Further adding craters with upper age limits (s^{up}) gives the *full* data sets. I do not make any cut on the age uncertainties. The data sets are as follows

basic150 (32 craters) age \leq 150 Myr, σ_t original

ext150 (36 craters) age \leq 150 Myr, σ_t original or assigned

full150 (48 craters) ext150 plus craters with $s^{\text{up}} \leq$ 150 Myr

basic250 (42 craters) age \leq 250 Myr, σ_t original

ext250 (46 craters) age \leq 250 Myr, σ_t original or assigned

full250 (59 craters) ext250 plus craters with $s^{\text{up}} \leq$ 250 Myr

large400 (18 craters) age \leq 400 Myr, $d >$ 35 km, σ_t original

This final data set, large400, extends further back in time, but now we definitely expect a bias of preferential preservation and discovery for larger craters above 5 km. Following previous studies (see section 6), I therefore just retain craters with much large diameters in an attempt to avoid this bias. In addition to the 14 craters matching craters from Table 1, the four additional (older) craters in this set are: Clearwater West (290 ± 20 Myr, 36 km); Charlevoix (342 ± 15 Myr, 54 km); Woodleigh (364 ± 8 Myr, 40 km); Siljan (376.8 ± 1.7 Myr, 52 km).

3 METHOD

I now introduce the time series analysis method, showing how to model a generic time-of-arrival data set with a probabilistic model and how to calculate its evidence.

3.1 Bayesian hypothesis testing

The goal of hypothesis testing is to identify which of a set of hypotheses is best supported by the data. More quantitatively, we would like to determine $P(M|D)$, the probability that a hypothesis or model M is true given a set of measured data D . Here D is the ages for a set of craters. M could be “uniform distribution” or “periodic distribution”, for example.

Perhaps surprisingly, orthodox (or frequentist) statistics lacks a general framework for this problem and offers instead a number of recipes based on defining some statistic. These normally involve calculating the value for that statistic (e.g. χ^2), and comparing it with the value which would be achieved by some “random” (noise) model. As discussed at some length in the literature, some of these techniques are inconsistent or misleading, even when we just have two alternative hypotheses (e.g. Berger & Sellke 1987, Kass & Raftery 1996, Berger 2003, Christensen 2005, Bailer-Jones 2009; see also section 6). The Bayesian approach, in contrast, is direct and




Figure 1. The 59 craters listed in Table 1, excluding the 13 craters with upper limits on their ages. Several craters have identical or very similar ages so for the sake of this plot I shifted some ages by up to 0.5 Myr in order to distinguish them.




Figure 2. The 59 craters listed in Table 1, including the 13 craters with upper limit ages (plotted with the “J” symbol). There are two upper identical points at 70 Myr, 6 km. The error bars are the age uncertainties and in some cases are smaller than the size of the points. Note that diameter is plotted on a logarithmic scale.

often turns out to be quite simple. It inevitably involves a number of numerical integrals, but these can be solved with computers. For more background on Bayesian techniques in general see Jeffreys (2000), Jaynes (2003), MacKay (2003) or Gregory (2005).

To calculate $P(M|D)$ for one particular model M_0 we use Bayes’ theorem

$$\begin{aligned}
 P(M_0|D) &= \frac{P(D|M_0)P(M_0)}{P(D)} \\
 &= \frac{P(D|M_0)P(M_0)}{\sum_{k=0}^K P(D|M_k)P(M_k)} \\
 &= \frac{1}{1 + \frac{\sum_{k=1}^K P(D|M_k)P(M_k)}{P(D|M_0)P(M_0)}}
 \end{aligned} \tag{1}$$

where $k = 0 \dots K$ represents all plausible models. If there are only two, M_0 and M_1 , this simplifies to

$$P(M_0|D) = \left[1 + \frac{P(D|M_1)P(M_1)}{P(D|M_0)P(M_0)} \right]^{-1}. \tag{2}$$

This follows because implausible models are – by definition – those with negligible model prior probabilities, $P(M) \ll 1$. $P(D|M)$ is called the *evidence* for model M (derived in the next section). If we assign the two models equal prior probabilities, then the evidence ratio alone determines the posterior probability, $P(M_0|D)$. This evidence ratio is called the *Bayes factor*

$$BF_{10} = \frac{P(D|M_1)}{P(D|M_0)} \tag{3}$$

of model 1 with respect to model 0. When $BF_{10} = 1$ the posterior probability is 0.5 for both models. When $BF_{10} \gg 1$ then

$P(M_0|D) \simeq 1/BF$, and when $BF_{10} \ll 1$ then $P(M_0|D) \simeq 1 - BF_{10}$. If we calculate Bayes factors greater than 10 or less than 0.1 then we can start to claim “significant” evidence for one model over the other (e.g. Kass & Raftery 1996, Jeffreys 2000). I shall use Bayes factors throughout this article to compare models.

Given the Bayes factors for all models relative to M_0 , the posterior probability of this model is then

$$P(M_0|D) = \left[1 + \sum_{k=1}^{K} BF_{k0} \frac{P(M_k)}{P(M_0)} \right]^{-1}. \quad (4)$$

One difficulty of the Bayesian approach is that in order to calculate this posterior probability one must specify all plausible models (in order to get the correct summation needed to normalize the probabilities). This is often not possible (other than for simple two-way hypotheses). Yet even when we cannot identify all models, Bayes factors remain a valid way of comparing the relative merits of a set of models, and thus identifying the best of these.

3.2 Measurement model for time-of-arrival data

The critical characteristic of the crater time series is that it is just a list of times (with uncertainties), without any corresponding quantity. This is unlike most other time series encountered in astrophysics, such as a light curves or radial velocity time series. (Some authors have used crater size or inferred impact energy as the “dependent variable” in a time series analysis (e.g. Yabushita 2004), but this is risky given the significant uncertainties in diameters.)

Consider a single event j , with measured age s_j and corresponding uncertainty σ_j . This measured age is an *estimate* of the (unknown) true age, t_j . We express this uncertainty probabilistically. Assuming a Gaussian distribution for the measurement, the probability of observing the measurement s_j given the true age is

$$P(s_j|\sigma_j, t_j) = \frac{1}{\sqrt{2\pi}\sigma_j} e^{-(s_j-t_j)^2/2\sigma_j^2} \quad (5)$$

which is normalized with respect to s_j .²

3.3 Stochastic time series models

The goal is to compare plausible models which could produce the observed impact crater time series. Given the astrophysical context, we do not expect the sequence of crater impacts to be deterministic. For example, when we postulate that the time series is “periodic”, we do not expect the events to have a strict spacing, even if the ages were measured arbitrarily accurately. An exact, intrinsic rhythm perturbed only by measurement errors seems highly implausible a priori. We should instead understand periodic to mean a periodically varying probability of an impact. This is a *stochastic* model. It is described by $P(t_j|\theta, M)$, the probability of getting an impact at time t_j given model M with parameters θ . A simple periodic model would be a sinusoid, described by the two parameters period and phase.

² The maximum entropy principle tells us that if we only have the mean and standard deviation of a quantity, then the least informative (most conservative) probability distribution for this quantity is the Gaussian. Of course, if the standard deviation is a significant fraction of the age, then a Gaussian has a significant probability mass at negative age. As this problem effects no more than five events listed in the table, I consider this as an adequate approximation.




Figure 3. Principle of the likelihood calculation (equation 7)

3.4 Bayesian evidence for time-of-arrival data

I now put the above considerations together to derive the Bayesian evidence, $P(D|M)$.

The probability of observing data s_j from model M with parameters θ is $P(s_j|\sigma_j, \theta, M)$, the *likelihood* for one event. The time series model predicts the true age of an event, which is unknown. Applying the rules of probability we marginalize over this to get

$$\begin{aligned} P(s_j|\sigma_j, \theta, M) &= \int_{t_j} P(s_j, t_j|\sigma_j, \theta, M) dt_j \\ &= \int_{t_j} P(s_j|\sigma_j, t_j, \theta, M) P(t_j|\sigma_j, \theta, M) dt_j \\ &= \int_{t_j} P(s_j|\sigma_j, t_j) P(t_j|\theta, M) dt_j. \end{aligned} \quad (6)$$

The last step follows from conditional independence: in the first term – the measurement model – once t_j is specified s_j becomes conditionally independent of θ and M ; in the second term – the time series model – t_j is independent of σ_j . As the data are fixed, we consider both terms as functions of t_j . Both must be properly normalized probability density functions.

If we have a set of J events for which the ages and uncertainties have been estimated independently of one another, then the probability of observing these data $D = \{s_j\}$, the *likelihood*, is

$$\begin{aligned} P(D|\sigma, \theta, M) &= \prod_j P(s_j|\sigma_j, \theta, M) \\ &= \prod_j \int_{t_j} P(s_j|\sigma_j, t_j) P(t_j|\theta, M) dt_j. \end{aligned} \quad (7)$$

where $\sigma = \{\sigma_j\}$. The principle of this calculation is illustrated in Fig. 3: the likelihood of an event for a given model is the integral of the probability distribution for the event (eqn. 5) over the time series model, $P(t_j|\theta, M)$. Specific cases for the latter are introduced in section 3.6 below.

The *evidence* is obtained by marginalizing the likelihood over the parameter prior probability distribution, $P(\theta|M)$,

$$\begin{aligned} P(D|\sigma, M) &= \int_{\theta} P(D, \theta|\sigma, M) d\theta \\ &= \int_{\theta} P(D|\sigma, \theta, M) P(\theta|M) d\theta \end{aligned} \quad (8)$$

(where σ drops out of the second term due to conditional independence). For a given set of data (crater time series), we calculate this evidence for the different models we wish to compare, each parametrized by some parameters θ . The parameter prior $P(\theta|M)$ encapsulates our prior knowledge (i.e. independent of the data) of

Table 2. Stochastic time series models. Time t increases into the past. *SigProb:Neg* and *SigProb:Pos* are special cases of *SigProb* with $\lambda < 0$ and $\lambda > 0$ respectively. *SinProbX:Y* specifies *SinProb* for $X < T < Y$ (in Myr).

Name	$P_u(t \theta, M)$	parameters
UniProb	1	none
SinProb	$\frac{1}{2} (\cos[2\pi(t/T + \phi)] + 1)$	T, ϕ
SinBkgProb	$\frac{1}{2} (\cos[2\pi(t/T + \phi)] + 1) + b$	T, ϕ, b
SigProb	$(1 + e^{-(t-t_0)/\lambda})^{-1}$	λ, t_0
SinSigProb	SinProb + SigProb	T, ϕ, λ, t_0

the probabilities of different parameters, normally established from the context of the problem (see section 3.7). Note that the evidence is dependent on the measurement uncertainties (although I drop this explicit conditioning in the rest of this article).

A fundamental aspect of this evidence framework for model assessment is that we are not interested in the “optimal” value of the parameters for some model. We are not interested in “fitting” the model, but rather in determining how good the model is overall at explaining the data. The evidence measures this by averaging the likelihood, which is defined at fixed θ , over the prior for θ . The prior probability distribution is therefore important, and should be considered as part of the model definition. For example, a periodic model with a non-zero prior over the period range 10–50 Myr is distinct from a model with a non-zero prior over the range 50–100 Myr. The reason why we should integrate over the possible solutions rather than selecting the best one will be illustrated and discussed in some detail.

3.5 Inclusion of events with age limits

By representing the unknown, true age of an event probabilistically, we take into account the age uncertainties. It also allows us to include events which only have upper or lower limits on their ages (censored data).

Given a (hard) upper limit, s_j^{up} , all we know is that the true age is younger. A suitable – and uninformative – measurement model is to assume that the probability of measuring some value for the upper limit is constant for any age below the one we actually measured, s_j^{up} , and zero otherwise. In this case, we write equation 5 as

$$P(s_j^{\text{up}}|t_j) = \begin{cases} \frac{1}{s_j^{\text{up}}} & \text{when } 0 < t_j < s_j^{\text{up}} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

which is normalized. Recall that time increases into the past.

Events which have lower limits to their ages may be treated in the same way, provided we can also assign an upper limit to the possible age (required to normalize the probability density function). It is not obvious how to assign this. We might use the oldest event in the data set or the age of the oldest rocks searched for craters. As only two events with lower limits occur in the time ranged analysed (and both quite low, at 35 Myr), these are not included.

3.6 Impact cratering time series models

Table 2 lists the time series models which will be tested. The second column gives the (unnormalized) probability density function that an event occurs at time t for given parameters. (For given

parameters, these distributions must be normalized over the time span of the data when they are used in the likelihood integral.) *SinProb* is a sinusoid with period T and phase ϕ . *SinBkgProb* adds a constant background, b , to this. *SigProb* is a sigmoidal function parametrized by the steepness of the slope (λ) and the time of the centre of the slope (t_0). It is used to model a monotonic trend in the event probability, with $\lambda < 0$ giving a decrease in probability with (look back) time. *SinSigProb* models a periodic signal on top of a trend. Examples of these functions are shown later.

3.7 Choice of parametrization and parameter prior distribution

To calculate the evidence (equation 8) for the models, we integrate over the parameters. In order that this integral converge we must either adopt a proper prior or, equivalently, we must adopt a finite parameter range over which the model is defined. In the interests of keeping assumptions limited and the results intuitive, I adopt a uniform prior over a finite range for all parameters,

$$P_u(\theta|M) = \begin{cases} 1 & \text{when } \theta_{\min} < \theta < \theta_{\max} \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

The u subscript denotes that this is not normalized (divide by $\Delta\theta = \theta_{\max} - \theta_{\min}$ to normalize)³.

It is important to realise that the adopted parameter range is an intrinsic part of the model. Hence, rather than talking about the model *SinProb*, for example, we should talk about the model defined over some period range. For this reason I will refer to models like *SinProb10:50*, which means the *SinProb* model for the period range 10–50 Myr.

The evidence of course also depends on the shape of the prior, and I adopt a uniform prior mostly to represent ignorance. But uniform over what? It would be equally valid to parametrize the periodic models in terms of frequency ($\omega = 1/T$) and adopt a uniform prior over that, for example. The transformation between the priors is $P_T = P_\omega T^{-2}$, so a prior uniform in frequency is non-uniform in period. In particular, shorter periods will achieve more weight in the evidence calculation. To assess the impact of this choice, I have repeated all of the analyses described in sections 4 (simulations) and 5 (EID results) which use periodic models with a prior uniform in frequency. It turns out that this change makes bi significant difference to the results, and does not change the conclusions. (The issue of priors is discussed further in section 6.)

The ranges for the other parameters are as follows. The phase parameter is only defined in the range 0–1, so it is natural to always use this full range (no reason to exclude any phases). For the background parameter b in *SinBkgProb*, I examine the evidence over a range 0–5, the upper limit set by the intuition that with much larger backgrounds it will be hard to detect a periodic signal (this was later found to be the case). For the parameters of *SigProb*, I cover the range of λ from 0 to ± 100 . This encompasses models ranging from a step function ($\lambda = 0$) to a virtually flat function over the period of interest (see the figures in section 4). The range of t_0 is chosen as 0–275 in order to move the “crossing point” of the sigmoid across the whole basic250 time range.

³ Normalization is essential. It ensures that model complexity is taken into account by the model comparison. Note, therefore, that we are not modelling the absolute rate of impacts, just the probability.




Figure 4. Six examples of simulated time series drawn from the model *UniProb* (each with 42 events)

3.8 Numerical estimation

The integral in equation 8 is estimated numerically using

$$\int_x f(x)dx \approx \sum_{n=1}^{n=N} f(x_n)\delta x = \frac{\Delta X}{N} \sum_{n=1}^{n=N} f(x_n) \quad (11)$$

where $f(x)$ is any continuous integrable function, the sample $\{x_n\}$ is drawn from a uniform distribution, ΔX is the range of x from which they are drawn and $\delta x = \Delta X/N$ is the average spacing between the samples. (For equation 7 I just used a regular dense sampling.) $P(\theta|M) = P_u(\theta|M)/\Delta\theta = 1/\Delta\theta$ for all the priors I use, so the numerical integration simplifies to

$$P(D|M) \approx \frac{1}{N} \sum_{\theta=\theta_{\min}}^{\theta=\theta_{\max}} P(D|\theta, M) \quad (12)$$

with the samples drawn from a random uniform distribution. Thus with uniform priors, the evidence is just the likelihood averaged over the range of the parameters. This is evaluated at several million randomly selected points, more than enough to ensure a very high signal-to-noise ratio in the calculated evidence.

4 TESTS ON SIMULATED TIME SERIES

To test the method I apply it to numerous simulated data sets. The goal is to determine whether we can identify the underlying signature in the data, and with what significance, and whether we sometimes identify the wrong model. I will demonstrate, for example, that simply identifying a “significant” peak in the periodogram can lead us to the wrong conclusion.

I generate stochastic time series using the same probabilistic models described in section 3.6. For each model and set of parameters (e.g. period and phase in *SinProb*) I draw events independently at random from the probability distribution. I generate several random time series for each parameter combination. In all cases a time series comprises 42 events over a time span of 250 Myr – characteristic of the basic250 data set – and σ_t is unity for all events.

Figures 4, 5 and 6 show example time series drawn from *UniProb*, *SinProb* and *SigProb* respectively. Note how different the time series can be even when drawn from the same model with the same parameters. It is also interesting – but not surprising – that stretches of some of the uniform random distribution appear




Figure 5. Six examples of simulated time series drawn from the model *SinProb* (each with 42 events). The model is plotted in red, the randomly selected events in black (in a split panel to aid viewing of these stochastic time series). The model has a period of 50 Myr on the left and 100 Myr on the right (phase = 0.5 in both cases)




Figure 6. Six examples of simulated time series drawn from model *SigProb* (each with 42 events). The model is plotted in red, the randomly selected events in black (in a split panel to aid viewing of these stochastic time series). The model has $\lambda = -30$ on the left and $\lambda = -70$ on the right ($t_0 = 150$ in both cases)

almost periodic. Conversely, the time series drawn from *SinProb* do not always look very periodic. In all cases the distribution can be very uneven and/or show clustering and gaps. Searching for patterns by eye can be quite misleading.

I report here just the results of the periodic models using the uniform prior over period. The results of using a uniform prior over frequency are very similar (I mention a few in the footnotes). Sometimes the evidence is higher or lower, but it does not change the strength of any of the conclusions reported. The maximum likelihood solutions are, in virtually all simulations, identical in parameters and likelihood to within 0.1%.

4.1 Periodic data sets

I generate data sets from *SinProb* at several different fixed periods and phases. For each time series I calculate the evidence and Bayes factors (evidence ratios) for a number of models.

For the *SinProb* model, the likelihood is calculated at several million values of period and phase drawn from the uniform prior distribution for periods of 10–550 Myr and phases of 0–1. These likelihoods are plotted as a density plot in Fig. 7 for one particular simulated time series with true period 35 Myr and phase of 0.75 (shown in Fig. 8). There is large variation in the likelihood. The overall evidence for this model for the full period range 10–550 Myr is found by averaging the likelihoods, and is 8.12×10^{-101} . The evidence for *UniProb* is 8.89×10^{-102} , giving a Bayes factor of 9.1. This is not quite significant according to standard criteria, and would appear to suggest a lack of evidence for the periodic model at first. However, we have searched over a wide range of “periods”: up to twice the time span of the data. We see in Fig. 7 that these have very low likelihoods, bringing down the average (the evidence) for *SinProb*. So if we instead average over periods of 10–125 Myr then we get an evidence for this “properly periodic” model of 3.22×10^{-100} , a Bayes factor of 36 relative to *UniProb*. This is good evidence that the periodic model is the better of the two.⁴ This is also true for the model over intermediate periods: $\text{BF}(\text{SinProb}10:250/\text{UniProb}) = 24.5$.

We can also calculate the evidence over very narrow ranges of period (for all phases). If we do this for consecutive ranges, we get a *Bayesian periodogram*: the variation of likelihood with period (or frequency). This is equivalent to marginalizing the two-dimensional likelihood over phase. It is plotted for the present example in Fig. 9. As these likelihoods have uninterpretable absolute values, I divide (“normalize”) the likelihoods by the evidence of the *UniProb* model. The periodogram gives the Bayes factor of *SinProb* at a specific period (strictly, over a very narrow period range) relative to *UniProb*, which I denote as $\text{BF}(\text{SinProb}(t)/\text{UniProb})$.

We see a (very) significant peak at just a single period, 34.75 ± 0.6 Myr (the uncertainty being the half-width at half maximum, HWHM). As we have already established evidence for the periodic model overall, this indicates that the periodic model at the detected period is a good explanation for these data.⁵

The maximum likelihood solution is at a period of 34.5 Myr and a phase of 0.73. Its likelihood is 59 000 times higher than the evidence for *UniProb*. This is larger than the peak in the periodogram




Figure 7. Likelihood distribution as a function of period and phase for the *SinProb* model calculated on the simulated data set shown in Fig. 8 (likelihoods were calculated up to 550 Myr, but are only plotted up to 250 Myr.). The logarithm (base 10) of the likelihoods are shown on a colour scale: it spans 35 orders of magnitude. White regions are those where the likelihood drops below the minimum plotted.




Figure 8. A simulated time series (black lines) drawn from *SinProb* with $T = 35$ Myr and $\phi = 0.75$ (red curve)




Figure 9. Bayesian periodogram obtained by marginalizing the 2D likelihood distribution in Fig. 7 over the phase. The inset shows a zoom on the 10–50 Myr portion

⁴ The evidence for *SinProb*:10:125 with the uniform frequency prior is 3.67×10^{-100} , hardly any different.

⁵ The peak in the periodogram for model with the uniform frequency prior is at 34.64 Myr, and has a BF 86% as large.

because now we have also found the optimum phase. However, because this model has two parameters which have been fit to the data, this is not a representative measure of the “significance” of a periodic model in which the parameters are not known a priori. This will be discussed further below.

This example is actually one in which the simulated time series is relatively non-periodic. Many of the random generations produce much more periodic data sets which achieve peaks in the periodogram of 10^5 or more and a $\text{BF}(\text{SinProb}10:125/\text{UniProb})$ of hundreds to thousands.

The method is very efficient at finding periods in these data sets. Of 100 different random time series drawn from *SinProb* with period 35 Myr and phase 0.5, all 100 showed a (very) significant Bayes factor (relative to *UniProb*) for both the overall model (*SinProb*10:125) and at the peak period (which agreed with true period in 99 cases).

I obtain very similar results for numerous other time series simulated at other periods and phases: the evidence for the *SinProb*10:125 model is almost always very significant, and there is also always a very strong peak in the Bayesian periodogram. (As one would expect the peak is wider – the period less certain – for longer periods.) The method can also detect periods between 125 and 250 Myr, although periods longer than 250 Myr (where there is less than a complete cycle in the data) are sometimes not detected. Periods down to the lower limit searched for, 10 Myr, can also be recovered reliably.

It must be stressed that the above has only established evidence of *SinProb* relative to *UniProb*. In Bayesian hypothesis testing we can only ever *compare* models. I therefore calculated the evidence for the trend model, *SigProb*, defined over the parameter ranges $-100 \leq \lambda \leq +100$, $0 < t_0 < 275$. For the above example (Fig. 8), we get a Bayes factor $\text{BF}(\text{SinProb}10:125/\text{SigProb}) = 155$, clearly favouring the periodic model. This is also seen for other simulated time series at the same and other periods: the Bayes factor $\text{BF}(\text{SinProb}/\text{SigProb})$ is typically 10–1000 for input periods below 250 Myr. For longer periods the *SigProb* model sometimes dominates, depending on the exact data set. This is because such long “periods” are really trends, and these may be fit better with *SigProb*.

In conclusion, if the time series is periodic (in the sense of a periodic event probability), then the method strongly favours the periodic model over both a uniform one and a monotonic trend model. We conclude this primarily from the significant (large) Bayes factor for the model as a whole (i.e. a broad period range for all phases), and then from a significant peak in the Bayesian periodogram (the Bayes factor over a very narrow period range).

4.2 Uniform and trend data sets

Let us now examine whether we erroneously favour the *SinProb* model – and possibly detect artificial periods – in non-periodic data.

Uniform data sets

For this purpose I simulated 50 stochastic time series from the *UniProb* model and calculated the evidence for both *UniProb* and *SinProb*. In none of these 50 cases is $\text{BF}(\text{SinProb}10:125/\text{UniProb})$ significant: most of the time series give 10^{-2} to 10^{-3} (the full range is 3×10^{-4} to 1.1). $\text{BF}(\text{SinProb}10:250/\text{UniProb})$ is likewise much less than one. We would therefore correctly conclude that a uniform random distribution is a much better model than a




Figure 10. A simulated time series (black lines) drawn from *SigProb* with $\lambda = -100$ Myr and $t_0 = 0$ Myr (red curve)

periodic one. (If no other model were plausible, we could further conclude with some confidence that this is the “best” model.)

However, in 10 of the cases we nonetheless observe a significant peak in the periodogram at periods less than 125 Myr, where significant means $\text{BF}(\text{SinProb}(t)/\text{UniProb}) > 10$. This is because even a uniform random time series with 42 points can happen to show a weak periodicity at *some* period, even though the periodic model overall is a poorer explanation for the data than the uniform one.⁶ If we took just the identification of these “significant” peaks to be evidence for a period (as is done in frequentist periodogram analyses), we would make a false positive claim in 20% of cases. This underlines the importance of basing a conclusion on the evidence for the model as a whole, rather than a specific fit.

Trend data sets

We see the same general phenomena among a set of time series generated from *SigProb*: The evidence for *SinProb*10:125 is much lower than the evidence for *SigProb*, leading us to correctly conclude a lack of evidence for periodicity. But once again there are spurious “significant” peaks in the periodogram. The reason for this is that although *SinProb* sometimes gives high evidence at a very specific period, once averaged over a broader period the evidence is much less. If the peak had been much higher (as was the case for the truly periodic data sets above, e.g. Fig. 9), then averaging over a broader period reduces the evidence, but not enough to make it insignificant. The model evidence is a balance between the quality of a specific fit, and how much of the parameter space produces a good fit. (This is discussed further in section 6, where it can be understood in terms of Occam factors). As we had no prior reason to suspect a periodicity *at the peak found*, it is misguided to focus only on that peak – and that may anyway give the wrong conclusion, as we have just seen.

An example time series drawn from the trend model is shown in Fig. 10. The likelihood distribution for the *SigProb* model is shown in Fig. 11, with a colour scale spanning 15 orders of magnitude. Positive trends (probability increasing with look back time) are heavily disfavoured. In contrast, a very broad range of parameter space for $\lambda < 0$ has large evidence: the model is not very sensitive to the exact parameters. This contrasts with the simulated periodic data sets, where the evidence for the overall model came from a very narrow range of the parameters with very high likelihood (see Figs 7 and 9).

⁶ Half of the time series in fact show a peak in the periodogram with $\text{BF}(\text{SinProb}(t)/\text{UniProb})$ in the range 1 to 10. We clearly must resist the temptation to read too much into such low significance peaks!




Figure 11. Log likelihood distribution as a function of λ and t_0 in the *SigProb* model calculated for the simulated data set shown in Fig. 10. The $\log(\text{Evidence})$ for the overall model is -97.8 (and 0.3 higher for *SigProb:Neg*). For comparison, $\log(\text{Evidence}) = -100.9$ for *UniProb*.

It is interesting to note that the evidence for *SinProb10:125* on these trend data sets is low even when compared to *UniProb*: The Bayes factors range from 0.7 to 5×10^{-5} in 39 of 40 simulations (the other value was 2). So even when comparing against an overly simple model, we get no evidence for periodicity. Yet the in the *periodograms* normalized to the evidence in *UniProb* we see apparently significant peaks, i.e. with $\text{BF}(\text{SinProb}(t)/\text{UniProb}) > 10$. These peaks are irrelevant, however, not only because *SinProb* as a whole is disfavoured but also because *UniProb* is a poor model for the data. That is, we can increase the *apparent* significance of a model (here *SinProb*) by comparing it with an overly-simple alternative model (here *UniProb*), a common mistake in hypothesis testing.

4.3 Compound data sets

It is plausible that the impact cratering phenomenon comprises both a periodic and a non-periodic component (e.g. Grieve et al. 1988, Lyytinen et al. 2009). It is generally more difficult to identify evidence for more complex models given a limited amount of data and the stochastic nature of the phenomenon.

This difficulty is confirmed by simulations. I simulated 48 time series drawn from *SinBkgProb* for several periods between 10 and 250 Myr and $b = 1$, i.e. equal amplitudes of the periodic and uniform components (Table 2). In many cases the evidences for *SinBkgProb10:550*, *SinBkgProb10:125* and *UniProb* are about equal, so the true model is not favoured over a pure “background” model. (*SigProb* achieves a very low evidence in comparison: we do not erroneously claim a trend when there is not one.)

We see much the same for several time series drawn from *SinSigProb* at each of three different periods (20, 35, 50 Myr) with a fixed trend ($\lambda = -60$, $t_0 = 100$). The Bayes factors of the true models relative to just the trend models, *SigProb*, lie between 0.5

and 7.⁷ Thus although the true model is not explicitly disfavoured, these values are not high enough to claim significant evidence to favour it, and when lacking evidence in favour of a more complex model, we would probably prefer the simpler one.

It seems that the method is conservative, having difficulties to identify a period in a stochastic time series which includes a large (here 50% amplitude) non-periodic component. A similar conclusion was reached by Lyytinen et al. (2009) using a different approach. However, I have only performed a few tests with these compound data sets, and only for a very narrow part of the input parameter space. As the sensitivity is likely to vary over the input parameter space, more extensive tests are necessary.

In conclusion of this section on simulations, I have found that if the time series is a stochastic one drawn from a model with a periodic distribution, then there is both significant evidence for periodicity and we can identify the true period. Conversely, if data are drawn from a uniform or trend distribution, we find significant evidence for the correct model, and do not erroneously identify periodicity. Preliminary tests indicate difficulty in favouring more complex compound models (when true) over a simpler one. It may be that 42 events is simply inadequate to support more complex models.

5 RESULTS

Armed with the experience of how the models respond to time series of known origin, I now turn to analysing the real cratering data. For quick reference, the results are summarized in Table 3. For the periodic models I only report results using a prior uniform over period, because the results for using a prior uniform over frequency are the same in all relevant respects. An example is shown in the appendix.

I should point out that although I calculate likelihoods for the periods models for periods up to twice the time span, I do this only to illustrate that long “periods” can only really be interpreted as long-term trends, and so are better modelled by the *SigProb* model. The main results for “periodicity” are for the evidence calculated over narrower period ranges.

5.1 Data sets: basic150 and ext150

I start with the basic150 data set, the 32 craters over the past 150 Myr. The likelihood distribution for the *SinProb* model is shown in Fig. 12. The evidence for *SinProb10:300* is 9.64×10^{-72} . This compares to 1.63×10^{-70} for the *UniProb* model, giving a Bayes factor of 0.060, or significant evidence in favour of *UniProb*. The maximum likelihood over this parameter range occurs at (period, phase) = (11.7, 0.73) with a likelihood of 8.00×10^{-69} , or $\text{BF}(\text{SinProb}(t)/\text{UniProb}) = 49$. Although this tells us that this very specific fit explains the data better than a uniform random distribution, it only comes about because we have tuned two parameters (period and phase). *UniProb* has no free parameters and cannot be tuned, so a simple comparison of maximum likelihood does not take into account the model complexity. This is why we must use the evidence for the models as a whole. After all, we wanted to look for evidence of periodicity in general, not for evidence of “a period of 11.7 Myr and phase 0.73”, and more complex models could be defined which for which there are even higher likelihood fits.

⁷ When using a uniform prior over frequency the Bayes factors come out to within about 20% of the same values.

Table 3. The evidence (equation 8) for different models and data sets. The Bayes factor is the ratio of two evidences (for a given data set). See the text for the exact parameter ranges used in each case.

	basic150	full150	basic250	ext250	full250	large400
UniProb	$1.63e-70$	$2.30e-105$	$1.03e-103$	$1.68e-113$	$3.27e-145$	$1.25e-47$
SinProb10:50	$2.67e-71$	$1.41e-105$	$2.87e-104$	$1.11e-113$	$4.83e-145$	$8.89e-48$
SinProb10:125	$1.07e-71$	$5.56e-106$	$1.15e-104$	$4.18e-114$	$1.71e-145$	$1.01e-47$
SinProb10:300	$9.64e-72$					
SinProb10:400						$3.35e-48$
SinProb10:550			$8.57e-103$			
SinBkgProb10:50	$1.62e-70$		$1.11e-103$	$2.23e-113$	$5.22e-145$	
SinBkgProb10:125	$1.58e-70$		$9.71e-104$	$1.79e-113$	$3.91e-145$	
SinBkgProb10:300	$1.74e-70$					
SigProb	$1.35e-70$	$1.28e-103$	$8.64e-102$	$1.21e-110$	$7.66e-139$	$2.93e-48$
SigProb:Neg	$2.64e-70$	$2.56e-103$	$1.73e-101$	$2.41e-110$	$1.53e-138$	$4.83e-48$
SigProb:Pos	$5.70e-72$	$3.73e-108$	$1.47e-106$	$8.72e-117$	$3.01e-149$	$1.03e-48$
SinSigProb			$6.94e-102$	$6.81e-111$	$7.97e-140$	




Figure 12. Log likelihood distribution as a function of period and phase for the *SinProb* model for the basic150 data set

Limiting the evidence calculation to shorter periods of 10–50 Myr, we get a Bayes factor of *SinProb10:50* relative to *UniProb* of 0.16, still insignificant evidence for periodicities.

The Bayesian periodogram (section 4.1) is shown in Fig. 13. There is no significant evidence for periodicity at any period. Recall that, in the simulations, we obtained a peak (at the true period) with far higher Bayes factors (significance) when the data were drawn from a sinusoidal probability distribution.

In the appendix I repeat this analysis using a prior uniform in frequency for *SinProb*.

If the impact history comprises both constant and periodic components, then a model which reflects this may identify a periodicity better. To examine this I calculate likelihoods for *SinBkgProb*, which adds a variable constant background term, b , to *SinProb*. For the full period range and for $0 \leq b \leq 5$ the evidence is 1.74×10^{-70} , a Bayes factor relative to *UniProb* of 1.1. If we limit the period range to 10–50 Myr this rises only to 1.2. The pe-




Figure 13. Bayesian periodogram for the *SinProb* model for the basic150 data set

riodogram shows smaller Bayes factors than in Fig. 13, even when calculated for limited ranges of b . Hence *SinBkgProb* describes the data no better than a purely uniform random distribution, which we should arguably then prefer (see section 6).

So far the uniform random distribution describes the data as well as or better than a periodic one. This does not mean that this is *the* best model, however: we can only assess what we explicitly test. To look for the evidence for a trend in the data, I calculate the evidence for the *SigProb* over the parameter range $-100 \leq \lambda \leq +100$, $0 < t_0 < 150$. This gives 1.35×10^{-70} , a Bayes factor relative to *UniProb* of 0.83. Splitting *SigProb* into two distinct models, one for $\lambda < 0$ (*SigProb:Neg*) and the other for $\lambda > 0$ (*SigProb:Pos*), the Bayes factors are 1.63 and 0.035 respectively. The latter – an increase in impact probability with look back time – is disfavoured.

Performing the same analyses on the ext150 data set gives very similar results: Adding these four events does not make the evidence for periodicity nor for a trend significant.

In summary, we have no evidence for periodicity in the basic150 or ext150 data sets. Of the models tested, both *UniProb* and *SigProb:Neg* are more or less equally probable explanations.

Given a lack of strong evidence in favour of the more complex trend model, I conclude that the simpler, uniform random distribution is an adequate – and plausible – explanation for impact craters over the past 150 Myr.

5.2 Data set: full150

I now add the 12 craters which have upper age limits below 150 Myr, using the approach explained in section 3.5. The evidence for the models tested (same parameter ranges for basic150) are listed in Table 3. There is now significantly more evidence for the trend model than for the uniform one. More specifically, the negative trend (*SigProb:Neg*) is hugely favoured over the positive one.⁸ As the new data are *upper* age limits, it is not surprising that their inclusion increases the evidence for *SigProb:Neg*, although the evidence is very strong: $\text{BF}(\text{SigProb:Neg}/\text{UniProb}) = 111$.

Thus adding the 12 craters with upper age limits in a conservative manner (a flat probability distribution) has made *SigProb:Neg* much more probable than *UniProb*. *SinProb* remains unlikely.

Let us now extend the data set back to 250 Myr BP.

5.3 Data set: basic250

This data set comprises 42 events. The evidence for *UniProb* is 1.03×10^{-103} compared to 8.57×10^{-103} for *SinProb10:550*, a Bayes factor of 8.3. The top-left panel of Fig. 14 shows the likelihood distribution. The largest likelihoods are in the upper right quadrant, for periods greater than 200 Myr and phases above 0.5. Indeed, the maximum likelihood solution is at a period of 550 Myr with a phase of 0.94: it is plotted in the top-right panel of the same figure. This “period” is twice the duration of the data and is actually modelling a trend of decreasing probability with look back time. Marginalizing the likelihood distribution over all phases to produce the periodogram (Fig. 15), we see that all of the significant periods are at these very long trend periods. The evidence for true periods – *SinProb10:125* – is 1.15×10^{-104} , a Bayes factor relative to *UniProb* of 0.11.

Given that the likelihood varies considerably across the parameter space, it is informative to examine the model at various parameter “solutions”. Five examples are shown in Fig. 14. The two panels in the central row have the same period but different phases: the right-hand one gives an *increasing* probability of events with look back time and is hugely disfavoured by the data compared to the decreasing trend. The lower two panels are local maxima in the likelihood distribution at shorter periods. They look as though they could reasonably produce the observed data, bearing in mind that these are stochastic models. However, even the better of the two has a likelihood 10–100 times smaller than the longer period “trend” solutions.

Introducing a constant background into the model (*SinBkg-Prob*), we again find that the highest likelihood solutions (and the only ones more likely than *UniProb*) are at long “periods”. Solutions with $b > 0.5$ are favoured compared to those with a smaller (or zero) background. There is no evidence for any proper period (< 125 Myr), with Bayes factors relative to *UniProb* of no more than 0.94, even if we examine narrow ranges of b . So there is no evidence for periodicity superimposed on a constant background

probability (although recall from section 4.3 that it may be hard to identify such a model with confidence).

Given the above evidence for a trend, I fit *SigProb* over the parameter range $-100 \leq \lambda \leq +100$, $0 < t_0 < 275$. The overall evidence is 8.64×10^{-102} . This includes both positive and negative trends: Fig. 16 shows that the former ($\lambda > 0$) have much lower likelihoods. Splitting this model into two, we find that the evidence for $\lambda < 0$ (*SigProb:Neg*) is 1.73×10^{-101} , a Bayes factor of 167 relative to *UniProb*.

As *SigProb:Neg* is much more plausible than *UniProb*, this tells us that *UniProb* is an inappropriate reference model for the periodogram in Fig. 15. By comparing the evidence at each period to a model which predicts the data poorly, the periodic model may of course look good in comparison. But this is not the same as saying that the period is significant, because other (plausible) models may predict the data even better, as is the case here. This mistake is often made when assessing the significance of the classical (e.g. Lomb–Scargle) periodogram. One must be careful to choose an appropriate “background” or “reference” model for comparison. When using the evidence for *SigProb:Neg* as the reference in the periodogram, the Bayes factors in Fig. 15 are all reduced by a factor of 167, rendering all peaks – and even the long “trend” periods – entirely insignificant. Overall, the negative trend model is far more probable than the periodic one: $\text{BF}(\text{SigProb:Neg}/\text{SinProb10:125}) = 1500$.

Not only is *SigProb:Neg* the favoured model, it also has a large likelihood over a wide range of t_0 and λ . In other words, the evidence is not very sensitive to the exact parameter settings (nor to the prior). This can be seen in Fig. 17: half of the parameter space below $\lambda = 0$ has a likelihood within a factor of 10 of the maximum. (The model has a large Occam factor: see section 6) The maximum likelihood solution is plotted over the data in Fig. 18. Because the likelihood peak is so broad, we should not (need not) attribute much significance to this specific solution.

Is there evidence for a periodicity on top of this trend? We examine this using the four-parameter model *SinSigProb* using the same parameter range adopted for the periodic and trend models (namely all phases, periods of 10–125 Myr, $-100 \leq \lambda < 0$ and $0 < t_0 < 275$). The evidence is 6.94×10^{-102} , a Bayes factor with respect to *SigProb:Neg* of just 0.40. If we reduce the parameter range, e.g. to 10–50 Myr, $\lambda < 25$ and $t_0 < 200$, then the evidence hardly changes. We could possibly identify narrow, isolated regions of parameter space where *SinSigProb* is more probable, but this has no justification.

In summary of the basic250 data set, we have found significant evidence for a trend of decreasing probability of cratering with lookback time relative to both constant and periodic models. This is about the simplest model we can conceive after the constant one, and it also has more evidence than a model with a periodicity superimposed on either a trend or a constant background probability.

5.4 Data set: ext250

If I add the four craters for which I assigned age uncertainties, the evidence for the trend model relative to both the uniform model and the periodic models increases significantly (see Table 3; the parameters for *SinSigProb* are the same as used for basic250). This is not surprising because the additional craters are all comparatively young (1.07, 39, 46.5, 49 Myr).

⁸ As the evidence for *SigProb* is the average of the evidence for its two components, and one is five orders of magnitude smaller than the other, then the evidence for *SigProb:Neg* is just twice that of *SigProb*.

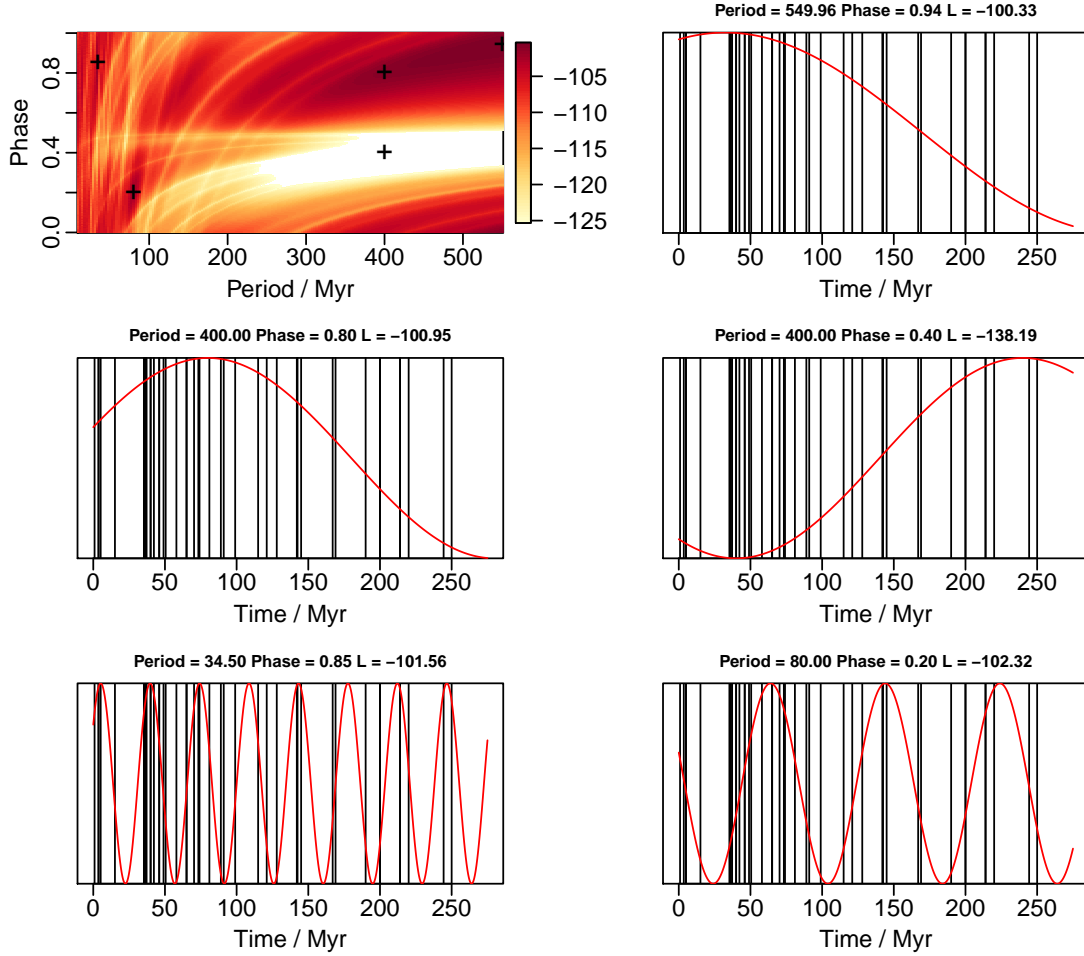


Figure 14. The top-left panel shows the log likelihood distribution as a function of period and phase for the *SinProb* model for the basic250 data set. The red curves in the five other panels show five possible “solutions” at five (T, ϕ) values corresponding to the five crosses plotted on the likelihood distribution. The black lines are the basic250 data. The log (base 10) likelihood and parameters of each solution are given at the top of each panel. The top-right panel is the maximum likelihood solution. The bottom left-panel is (close to) the maximum likelihood solution for short periods. For comparison, $\log(\text{Evidence})$ for *UniProb* and *SigProb:Neg* are -102.99 and -100.76 respectively.

5.5 Data set: full250

I now add the 13 craters with upper age limits, noting that 12 of them are below 150 Myr. As they are *upper* age limits, it is not surprising that their inclusion increases the evidence for *SigProb:Neg* further. But the increase is enormous: the evidence is 1.53×10^{-138} compared to 3.27×10^{-145} for *UniProb*, a Bayes factor of 4.7×10^6 .

The evidence for *SinProb10:125* is very small, giving a negligible Bayes factor relative to *SigProb:Neg*. The highest peak in the periodogram is still around 35 Myr, but its evidence is 10^5 times smaller than the evidence for *SigProb:Neg*. Modelling the data using a periodicity on top of a trend – with *SinSigProb* – increases the evidence for proper periods (i.e. below 125 Myr) dramatically over using just *SinProb*. But they are still insignificant compared to a pure trend. For example, the evidence for *SinSigProb* over periods of 10–50 Myr and $\lambda < 0$ is 7.97×10^{-140} , a Bayes factor of 0.052 relative to *SigProb:Neg*. All peaks in the periodogram are also insignificant. The data are better described with a pure trend.

In summary, adding the 13 craters with only upper age limits

radically increases the evidence for a negative trend, and radically decreases the evidence for either a periodicity or a periodicity plus negative trend, relative to the simple negative trend.

5.6 Data set: large400

The trend detected in the previous data sets might reflect a preservation bias in the geological record. When extending the analysis to older craters we can try to avoid this bias by only including larger craters. The large400 data set comprises 18 craters larger than 35 km with ages up to 400 Myr. The Bayes factor $\text{BF}(\text{SinProb10:400}/\text{UniProb})$ is 0.27. We see several low peaks in the periodogram, the four highest being 34 Myr (BF = 6.5), 18 Myr and 13.5 Myr (both BF = 5.5), and 100 Myr (BF = 4). As there is no prior reason to expect a period at any of these, we cannot simply select one and claim it as evidence (albeit marginal) for that period, especially given that there are so many low peaks (and only 18 data points). Again we must look at the overall evidence for periodicity. For the period range 10–50 Myr the Bayes factor




Figure 15. Bayesian periodogram for the *SinProb* model for the basic250 data set. If we normalize the periodogram against model *SigProb:Neg* instead of the *UniProb* then the Bayes factors are reduced by a factor of 167.




Figure 18. The maximum likelihood solution of *SigProb* (red curve) for the basic250 data set (black lines). It has $\lambda, t_0 = (-75, 98)$ and a likelihood of 6.04×10^{-101} .




Figure 16. Log likelihood distribution as a function of λ and t_0 for the *SigProb* model for the basic250 data set. The black cross marks the maximum likelihood solution.




Figure 17. As Fig. 16 but only for $\lambda < 0$ and with a shallower likelihood scale to better show the region around the peak.

$\text{BF}(\text{SinProb}10:50/\text{UniProb})$ is 0.71, implying that the periodic and uniform models describe the data equally well. But the periodic model with two free parameters for a pre-defined period range is arguably much less plausible a priori.

There is also no evidence for a positive or negative trend in the data, with Bayes factors below 0.4 for the *SigProb* models (see Table 3).

In summary, the large400 data set is most plausibly described by *UniProb*. Let us now turn to a discussion of the complete set of results as well as the method and principles on which it is based.

6 DISCUSSION

I have found significant evidence for a decrease in the cratering rate with lookback time (model *SigProb:Neg*) over the past 250 Myr for $d > 5$ km craters relative both to periodic models (*SinProb*, *SinBkgProb*, *SinSigProb*) and to a model with constant rate (*UniProb*). As there is no strong evidence for a trend in the past 150 Myr, this must come about primarily from a comparative lack of events between 150 and 250 Myr BP. This may now seem obvious from Figs. 1 and 2, but the analysis quantifies this and has also taken into account the age uncertainties. No such trend is found for larger craters ($d > 35$ km) over the past 400 Myr.

These results could be explained by a decreasing probability of preservation/discovery for older craters of size 5–35 km. However, studies of lunar cratering suggest that the cratering rate during the past 500 Myr was about twice as high as the average over the past 3.3 Gyr (e.g. Shoemaker 1983). More immediately relevant is the study of McEwen et al. (1997), who concluded that the cratering rate has increased up to the present by a factor of two during the past 300 Myr. If correct, and if the Earth is assumed to have experienced the same bombardment history, then this is consistent with my inferred increase of impact probability from 250 Myr BP up to the present. The single most probable solution for the trend model is shown in Fig. 18.

These conclusions are obviously based on the data we currently have. It is quite possible that a significant revision of the ages or the age uncertainties, or the inevitable discovery of more craters in the future, will lead us different conclusions based on

the same analysis. More craters may permit a better distinction between more complex models.

I will now discuss some aspects of the method, and compare the present analysis with previous work.

Significance assessment. The significance of a model can only be assessed relative to the significance of some other model. There is no absolute. In frequentist statistics one normally selects some “noise” or “background” model against which to compare a statistic measured on the real data. For example, with the classical periodogram the significance is usually determined from the distribution of the power achieved by a noise model. This may indicate that the periodic model is the better of the two, but both might be bad: there may be a third model which is better still. We saw an example of this in Fig. 14, where the bottom left-hand panel is the best-fit truly periodic solution. It was significant relative to *UniProb*, but insignificant relative to *SigProb:Neg*.

Why we should not rely solely on periodogram peaks. As has already been demonstrated in section 4, reliance on observing a peak in the periodogram – even when normalized to the true model – often results in erroneously claiming the periodic model to be a better explanation than the true one. The reason is that the periodogram has one free parameter (period), and we can sometimes find a specific value of this parameter which produces a better fit than the simpler uniform model (which has no free parameters). A model with even more free parameters may fit better still. But a model with *fitted* parameters is a priori less plausible than a model with no fitted parameters. Unless we have independent information to assign the model parameters, we cannot fit them and *then* compare that model on an equal footing with a model which has not been fit. Instead we must compare models “as a whole” (e.g. over some period range). We saw an example of this in section 4.2.

Occam factor. The conclusion of the previous discussion is not that more complex models are always penalized. They are not. What counts is how the plausibility of the model is changed in light of the data. This can be understood by the concept of the *Occam factor*. If the likelihood function is dominated by a single peak, then we can approximate the evidence (equation 8) with

$$\underbrace{P(D|M)}_{\text{Evidence}} = \underbrace{\mathcal{L}(\hat{\theta})}_{\text{best fit likelihood}} \times \underbrace{\frac{\Delta\theta_{\text{posterior}}}{\Delta\theta_{\text{prior}}}}_{\text{Occam factor}} \quad (13)$$

where $\mathcal{L}(\hat{\theta})$ is the likelihood at the best fit solution, $\Delta\theta_{\text{prior}}$ is the prior parameter range and $\Delta\theta_{\text{posterior}}$ is the posterior parameter range (the width of the likelihood peak) (see, e.g. MacKay 2003). The Occam factor (which is always less than or equal to one) measures the amount by which the plausible parameter volume shrinks on account of the data. For given $\mathcal{L}(\hat{\theta})$, a simple or general model will fit over a large part of the parameter space, so $\Delta\theta_{\text{prior}} \sim \Delta\theta_{\text{posterior}}$ and the Occam factor is not significantly less than one. We saw an example of this in Fig. 17. In contrast, a more complex model, or one which has to be more finely tuned to fit the data, will have a larger shrinkage, so $\Delta\theta_{\text{posterior}} \ll \Delta\theta_{\text{prior}}$. We saw this for the periodic models at short periods (e.g. Figs. 14 and 15), in which only a very specific period was a good fit to the data. In this case the Occam factor is small and the evidence is reduced. Of course, if the fit is good enough then $\mathcal{L}(\hat{\theta})$ will be large, perhaps large enough to dominate the Occam factor and to give the model a large evidence. We saw this with the simulated periodic time series for the *SinProb* model (Fig. 9).

This concept helps us to understand how the Bayesian approach accommodates model complexity, something generally lacking in frequentist approaches. If we assess a model’s evidence only by looking at the maximum likelihood solution (or the maximum over one parameter, the period), then we artificially compress the prior parameter range, increasing the Occam factor.

Parameter prior distributions. As the model evidence is the likelihood averaged over the prior parameter range (for uniform priors), this raises the issue of what this range should be. This is often the main perceived difficulty with Bayesian model comparison, and for some people this dependence on prior considerations is undesirable. Yet it is both logical and fundamentally unavoidable, because Bayesian or not, the prior parameter range is an intrinsic part of the model. Changing the parameter range changes the model, so will change the evidence. *SinProb10:50* is totally different from *SinProb100:150*, for example. If we are comfortable with deciding which are the plausible models to test, we must also be willing to decide what are the plausible parameter ranges to test. To some extent we can be guided by the context of the problem and the general properties of the data or the experiment, such as the sensitivity limits. For periodic models it seems obvious that we should use the whole phase range and that we should not include at “periods” much larger than the duration of observations (as these are more like trends). For *SigProb* I have actually used a rather broad range of its two parameters, even though some of this parameter space is a priori implausible, e.g. $\lambda = 0$ gives a probability of zero to one side of $t_0 = 0$.

More generally, the evidence is the likelihood averaged over the parameter prior distribution. There are often cases where we would not want to use a uniform distribution. It can be difficult to choose the “correct” prior distribution, and this choice may affect the results. Yet whether we like it or not, interpreting data is a subjective business: Just as we choose which experiments to perform, which data to ignore, and which models to test, so we must decide what model parameters are plausible. This seems preferable to ignoring prior knowledge or, worse, to pretending we are not using it.

In general, a probability density function is not invariant with respect to a nonlinear transformation of its parameters. As already discussed in section 3.7, I could equally well have used frequency rather than period to calculate the evidence for periodic models: there is no “natural” parameter here. This would not change the error model of the data (eqn. 5), and the value of $P(t_j|\theta, M)$ at period T is the same as when calculated at frequency $1/T$, so the likelihoods are unchanged. But as the model evidence is the average of the likelihoods over the prior, then the evidence would change if we adopted a prior which is uniform over frequency rather than over period. Thus the choice of parametrization becomes one of choice of prior. As neither parametrization is more natural than the other – a prior uniform over frequency does not seem to be more correct than one uniform over period – this remains a somewhat arbitrary choice. For this reason I repeated all of the analyses using periodic models with a prior uniform in frequency. Sometimes the evidence was slightly higher, sometimes lower, but the significance of the Bayes factors was not altered. The conclusions are robust to this change of prior/parametrization.

Model priors. I have used Bayes factors to compare pairs of models. Models are treated equally, so a significant deviation from unity gives evidence for one model over the other. However, if the models have different complexities (or rather, different prior plausibili-

ties) and the Bayes factor is unity, then rather than being equivocal we may tend to prefer the simpler model. This is because we normally give the less plausible model a lower model prior probability (equation 2). We can include these priors by reporting instead the *posterior odds*

$$\frac{P(D|M_k)}{P(D|M_0)} \frac{P(M_k)}{P(M_0)} = BF_{k0} \frac{P(M_k)}{P(M_0)}. \quad (14)$$

It is not obvious that all of the models I have considered should have equal model priors. For example, *SinSigProb* is arguably less plausible than *SigProb* a priori.

Posterior probabilities and plausible models. Ideally we would calculate model posterior probabilities, $P(M|D)$. This can only be done if we calculate the evidence for all *plausible* models, those with $P(M)$ which is not vanishingly small. In the current problem I have conceivably included most of the plausible parameter space for the stated models: These non-deterministic models are quite flexible in describing general shapes. For example, I found that a model with periodic Gaussians (with three parameters) looked very similar to *SinProb*. (Given sufficient physical reason, we could of course define more complex models, e.g. with a variable period or amplitude.) Assuming that *SigProb:Neg*, *SigProb:Pos*, *UniProb* and *SinProb10:125* are the only plausible models, and assigning them equal priors, then for the basic250 data set, the posterior probability of *SigProb:Neg* from equation 4 is $(1 + 8.50 \times 10^{-6} + 5.95 \times 10^{-3} + 6.65 \times 10^{-4})^{-1} = 0.993$.

Some problems with frequentist hypothesis testing. The motivation for the current work was to apply rigorous methodology to modelling impact crater time series, overcoming some limitations of frequentist hypothesis testing (for a discussion of these problems see, e.g., Berger & Sellke 1987, Marden 2000, Berger 2003, Jaynes 2003, Christensen 2005, Bailer-Jones 2009). To summarize, the main problems which can occur are as follows.

(i) Failure to take into account all plausible models. We see from equation 4 that the model posterior probability increases monotonically as the sum (over the alternative models) decreases. If we neglect plausible alternative models, the sum is smaller than it should be and the posterior probability of M_0 is artificially increased. This issue applies to any analysis, Bayesian or otherwise.

(ii) Incorrectly estimating significance by comparison to an inappropriate model (the “null” hypothesis).

(iii) Reliance on maximum likelihood or periodogram peak solutions without any regard to model plausibility/complexity (potentially leading to overfitting).

(iv) Failure to actually test the model of interest. A null hypothesis (M_0) is rejected and this is assumed to imply acceptance of the model of interest. This is only possible when there are only two plausible models which are mutually exclusive and exhaustive (rare in the physical sciences). Otherwise all models must be explicitly tested.

(v) Use of a statistic to reject a model if that statistic is *more extreme* than the value given by the data. This is the usual approach taken with p-values for example, and is used in most periodograms, including the Lomb–Scargle. The reference to values not observed cannot be justified, but the methods are forced to on account of the previous point (failure to actually calculate an evidence for the model of interest);

(vi) Incorrect interpretation of p-values. Frequentist statistics interprets a small value of $P(D|M_0)$ to be evidence for the alternative hypothesis, M_1 . But that is given by $P(M_1|D)$, which we

can *only* calculate if we also know $P(D|M_1)$ (equation 2). An example illustrates the difference. Suppose $P(D|M_0) = 0.01$. This is interpreted as evidence against the null hypothesis M_0 at $p = 0.01$. But if the evidence for the alternative hypothesis is higher, but still relatively low, e.g. $P(D|M_1) = 0.05$, then we see that $P(M_0|D) = 1/(1 + 0.05/0.01) = 0.17$ (for equal model priors). This is lower than $P(M_1|D)$, but is not low enough to “rule out M_0 at the 1% level”. P-values frequently overestimate the significance.

This is not to say that frequentist hypothesis testing and p-values are worthless. If we only have one model (e.g. “Gaussian random noise”), then it may be hard to define an explicit model for its complement, in which case Bayesian model comparison is awkward. Here a low p-value is a useful indication that this model may not be a good explanation, prompting the search for alternatives.

Previous studies of cratering periodicity. Using similar data to those used here, several studies have claimed evidence for periodicities in the impact crater data, whereas several others conclude no evidence for periodicity. (Some studies are quite old, when fewer impact craters had been discovered or dated). A non-exhaustive list of the studies and their main results follows.

- Alvarez & Muller (1984): period at 28.4 Myr for 11 craters with $t < 250$ Myr, $\sigma_t < 20$ Myr, $d > 5$ km using a classical periodogram method. Jetsu & Pelt (2000) claimed that this period (and potentially other claims of periodicity) is a artefact caused by rounding ages.

- Rampino & Stothers (1984): period at 31 Myr for 41 craters with $t < 250$ Myr using a correlation method. This work was strongly criticised by Stigler (1985).

- Grieve et al. (1985): periods found at 13.5, 18.5, 21 and 29 Myr for a set of 26 craters with $t < 250$ Myr, $d = 5$ –10 km using the method of Broadbent (1956). This method essentially defines a statistic which measures the deviation of a set of events from strict periodicity, and then estimates the probability that a uniform distribution would produce this value of the statistic (or smaller). No periodicity was found for $d > 10$ km. Partly because the extracted period depends on the data subset used, Grieve et al. concluded a lack of evidence for any true periodicity.

- Grieve et al. (1988): doubt cast on previously claimed periods around 30 Myr on a set of 27 craters, noting also that the correct period often cannot be found with the Broadbent method when there is a superimposed uniform random component. They found weak evidence for periods at 16 Myr and 18–20 Myr, the latter predominantly due to 10 craters in the past 40 Myr. (See also Grieve 1991.)

- Yabushita (1991): periods found at 16.5, 20 and 50 Myr for smaller craters ($d < 10$ km) in some data sets using the statistic of Broadbent, but different conclusions are reached depending on what significance test is adopted. He ultimately concludes that it is premature to claim evidence for periodicity.

- Yabushita (1996): a period at 30 Myr is found after examining various data sets (again using the Broadbent method), but it is found to be insignificant when a trend component (exponential decay) is included in the modelling.

- Montanari et al. (1998): no periodicity found for 33 craters with $t < 150$ Myr, $\sigma_t < 5$ Myr, $d > 5$ km using a clustering method which looks at the (uncertainty weighted) age differences between craters.

- Napier (1998): period at 13.4 Myr, suggested to be a harmonic of a 27 ± 1 Myr, from a set of 28 craters with $t < 250$ Myr, $\sigma_t < 10$ Myr, using a classical periodogram.

- Stothers (1998): period at 36 ± 1 Myr (using a variation of

the Broadbent method), but concluded to be insignificant. (He also criticises some earlier period searching work.)

- Yabushita (2004): period at 37.5 Myr from a set of 91 craters with $t < 400$ Myr (including craters which have upper age limits provided that limit is $t < 1$ Myr) using a Lomb–Scargle periodogram in which the crater size or impact energy is taken as the dependent variable. The analysis yields a p-value for periodicity between 0.02 and 0.1 (depending on which craters are used).

- Chang & Moon (2005): period at 26 Myr for various subsets with $t < 250$ Myr, $d > 5$ km using the Lomb–Scargle periodogram.

- Napier (2006): “weak” periods at 24, 35 and 42 Myr (depending on what we consider to be harmonics) in a set of 40 craters with $t < 250$ Myr, $\sigma_t < 10$ Myr, $d > 3$ km using a clustering method which examines the number of nearest neighbours.

Some studies have combined impact data with mass extinction time series (e.g. Napier 1998) or tried to identify correlations between them (e.g. Matsumoto & Kubotani 1996). Reviews and discussions can be found in Grieve (1991) and Grieve & Pesonen (1996). Numerical simulations of what signals can be detected in these kinds of data sets were presented in Heisler & Tremaine (1989) and Lyttinen et al. (2009).

The above summary makes clear that a large number of periods have been claimed, some of which can be identified as potential harmonics of others. Several of these periods, e.g. 11.5, 13.5, 18, 35 Myr, I also see as a local maximum in my likelihood distributions or periodograms. But in no case do I find any of them to be significant.

Most of the above studies employ frequentist hypothesis testing and suffer from one or more of the problems outlined previously. In particular, many compare the value of some statistic measured on the data to the values obtained by a noisy “reference” model. This is typically a uniform random distribution (akin to *UniProb*). Only the study of Yabushita (1996) examined a trend component in the analysis, and when this was included there was insufficient evidence for periodicity. As already demonstrated and discussed, *if the reference model is inappropriate and if other plausible models are ignored, then the significance of any periodicity is overestimated*. Moreover, by only considering the evidence in a single period we overfit the model, leading to a claim of periodicity where none exists. Note also that focusing on a single period (or narrow range) because it has been found in a previous study does not give an independent claim for periodicity, because we only have one crater record. This would amount to reusing the data, thereby increasing the evidence for the period artificially.

There may be a human desire to find and report periods. Several studies give some prominence to detected periods, but draw less attention to their limited statistical significance (e.g. Yabushita 1991, Stothers 1998). The motivation to do this may be a lack of confidence in the robustness of the test, anticipation that a later analysis will confirm the period(s) with significance, or because the period is close to another period (itself possibly insignificant) found in other studies, in biodiversity data, or in models of solar motion (e.g. Stothers 1998). But significance is everything: non-periodic models can give rise to superficially significant periods (section 4.2; see also Stigler & Wagner 1987). The (lack of) evidence for a periodicity in geological data or for an expected periodicity in astronomical phenomena is reviewed in Bailer-Jones (2009).

Summary of the main features of the method. The analysis method developed in this article is quite general and is not limited to analysis of impact crater time series. Its main features are

- operation on time-of-arrival data
- description of time series as stochastic models (more appropriate to the impact phenomena than deterministic models)
- consistent use of age uncertainties (obviating the need to remove “poor” data)
- ability to include craters with upper age limits (censored data) consistently
- use of Bayesian evidence calculation. Avoidance of p-values or ad hoc statistics
- comparison of multiple models (rather than relying on a single “null” hypothesis)
- use of proper parameter prior distributions, which are considered as an intrinsic part of the model.

The method should not be very sensitive to age errors provided the age uncertainties are approximately correct, although detailed, systematic testing of this has not yet been performed. We may also want to see how robust the conclusions are to the inclusion/removal of craters around the 5 km diameter limit. On the other hand, we have seen that the addition/removal of a few craters does not change the conclusions, as we would expect.

7 CONCLUSIONS

I find no evidence for a periodic variation in the impact crater rate the past 150 Myr or 250 Myr for craters with diameter above 5 km, relative to two other plausible – but quite broad – models: constant and monotonically varying probability with time. Compared to the uniform model, there is significant evidence for a monotonic decrease in the impact rate with look back time over the past 250 Myr, but not over the past 150 Myr. However, introducing craters with upper age limits into the analysis does give significant evidence for this trend even within the past 150 Myr. The physical interpretation is either an intrinsic variation in the impact probability or a variation in the crater preservation/discovery probability (i.e. we are less likely to find older craters). The former is consistent with some studies of lunar cratering. For very large craters ($d > 35$ km) over the past 400 Myr the best fitting model is the constant probability one, which is consistent with no preservation/discovery bias for such large craters. Given what we know about crater erosion and infilling, the preservation/discovery bias is a plausible explanation (e.g. Grieve 1991). It remains possible that there is a periodic variation on top of this. I find no evidence for this, although such a complex signal would be difficult to distinguish. Further simulations are necessary to explore what other kinds of signal could be reliably detected in these geological data.

Contrary to claims made in the literature, we can draw useful conclusions from events with large or variable age uncertainties, provided we use these uncertainties correctly. Larger uncertainties may make it harder to distinguish between models, but they do not invalidate the concept of model comparison.

Other studies claiming periodicity have probably overestimated the significance of the detected periods. This can occur if the significance is assessed relative to a null hypothesis of a model which is poorly supported by the data (here the uniform model), rather than to other non-periodic models which may be much better supported (here the trend model). The wrong conclusion can also be reached if we rely on a periodogram peak or the maximum

likelihood solution. I have shown via simulations how these can often claim a periodicity where none exists. The reason is that they fail to assess the evidence for the model as a whole, examining only the likelihood of an “overfit” solution.

ACKNOWLEDGEMENTS

I would like to thank Rainer Klement, Chao Liu, Kester Smith and Vivi Tsalmantza for fruitful discussions during the development of this work, as well as David Hogg and the referee for thoughtful comments on the manuscript. I am grateful to the staff of the Planetary and Space Science Centre at the University of New Brunswick, Canada, for their work in compiling and maintaining the Earth Impact Database, and to those who have contributed their data to it.

REFERENCES

- Alvarez W., 2003, *Astrobiology* 3, 153
- Alvarez L.W., Alvarez W., Asaro F., Michel H.V., 1980, *Science* 208, 1095
- Alvarez W., Muller R.A., 1984, *Nature* 308, 718
- Bailer-Jones C.A.L., *International Journal of Astrobiology* 8, 213
- Bahcall J.N., Bahcall S., 1985, *Nature* 316, 706
- Berger J.O., 2003, *Statistical Science* 18, 1
- Berger J.O., Sellke T., 1987, *Journal of the American Statistical Association* 82, 112
- Broadbent S.R., 1956, *Biometrika* 43, 32
- Chang H.-Y., Moon H.-K., 2005, *PASJ* 57, 487
- Christensen R., 2005, *The American Statistician* 59, 121
- Davis M., Hut P., Muller R.A., 1984, *Nature* 308, 715
- Deutsch A., Schärer U., 1994, *Meteoritics* 29, 301
- Le Feuvre M., Wicczorek M.A., 2008, *Icarus* 197, 291
- Gregory, P., 2005, *Bayesian logical data analysis for the physical sciences*, Cambridge University Press
- Gardner E., Nurmi P., Flynn C., Mikkola S., 2011, *MNRAS* 411, 947
- Grieve R.A.F., 1991, *Meteoritics* 26, 175
- Grieve R.A.F., Pesonen L.J., 1996, *Earth, Moon & Planets* 72, 357
- Grieve R.A.F., Sharpton V.L., Goodacre A.K., Garvin J.B., 1985, *Earth and Planetary Science Letters* 76, 1
- Grieve R.A.F., Sharpton V.L., Rupert J.D., Goodacre A.K., 1988, *Proceedings of the 18th LPSC*, 375
- Hallam A., 2004, *Catastrophes and lesser calamities. The causes of mass extinctions*, Oxford University Press
- Heisler J., Tremaine S., 1989, *Icarus* 77, 213
- Jaynes E.T., 2003, *Probability theory. The logic of science*, Cambridge University Press
- Jeffreys H., 2000, *Theory of Probability*, Cambridge University Press
- Jetsu L., Pelt J., 2000, *A&A* 353, 409
- Kass R.E., Raftery A.E., 1996, *Journal of the American Statistical Association* 90, 773
- Lyytinen J., Jetsu L., Kajatkari P., Porceddu S., 2009, *A&A* 499, 601
- Mackay D.J.C., 2003, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press
- Matsumoto M., Kubotani H., 1996, *MNRAS* 282, 1407
- Marden J.I., 2000, *Journal of the American Statistical Association* 95, 1316
- McEwen A.S., More J.M., Shoemaker E.M., 1997, *J. Geophysical Research* 102, 9231
- Montanari A., Campo Bagatin A., Farinella P., 1998, *Planet. Space Sci.* 46, 271
- Napier W.N., 1998, *Celestial Mechanics and Dynamical Astronomy* 69, 59
- Napier W.N., 2006, *MNRAS* 366, 977
- Rampino M.R., Stothers R.B., 1984, *Nature* 308, 709
- Shoemaker E.M., 1983, *Ann. Rev. Earth Planet. Science* 11, 461
- Shuter W.L.H., Klatt C., 1986, *ApJ* 301, 471
- Stigler S.M., 1985, *Nature* 313, 159
- Stigler S.M., Wagner M.J., 1987, *Science* 238, 940
- Stothers 1998, *MNRAS* 300, 1098
- Torbett M.V., Smoluchowski R., 1984, *Nature* 311, 641
- Wickramasinghe J.T., Napier W.M., 2008, *MNRAS* 387, 153
- Yabushita S., 1991, *MNRAS* 250, 481
- Yabushita S., 1996, *MNRAS* 279, 727
- Yabushita S., 2004, *MNRAS* 355, 51

APPENDIX A: MODELS WITH A UNIFORM FREQUENCY PRIOR

The periodic models described in section 3.6 have been re-run using frequency instead of period as the model parameter, and with a uniform prior over frequency. As discussed in the main paper, it turns out that this has no relevant impact on the results, i.e. the conclusions are robust with respect to this reparametrization or change of prior. So in the interests of brevity the results are shown here for just one model and data set.

Fig. A1 shows the likelihood distribution for the *SinProb* model applied to the basic150 data set. The corresponding periodogram, formed by marginalizing over the phase, is shown in Fig. A2. These two figures may be compared to the two corresponding figures for the model with the prior uniform over period, Figs. 12 and 13. The two periodograms actually show peaks at identical periods. Indeed, if we replot Fig. A2 in terms of period (i.e. transform the abscissa), we get a plot almost identical to Fig. 13, but with slightly higher values on the ordinate. The higher values are simply a result of the smoothing scale used to produce the plot: we get these slightly higher values in Fig. 13 too if we use a smaller smoothing scale. (This also reflects the fact that a fixed smoothing scale parameter in frequency corresponds to a variable one in period, and vice versa.) The maximum likelihood solution is exactly the same. The evidence for common period ranges is actually slightly different, e.g. 6.01×10^{-71} for periods of 10–50 Myr with the uniform frequency prior (Bayes factor of 0.37 with respect to *UniProb*), compared to 2.67×10^{-71} before. This reflects the different “weighting” of the likelihoods that comes with using a different prior (see equation 8). However, these changes are comparatively small, and do not alter the conclusions. In many other cases the differences are smaller.




Figure A1. Log likelihood distribution as a function of frequency and phase for the *SinProb* model with a prior uniform over frequency for the basic150 data set




Figure A2. Bayesian periodogram for the *SinProb* model with a prior uniform over frequency, for the basic150 data set